

POWER DOWN PROTOCOL FOR INTEGRATED CIRCUITS

Field of the Invention

The present invention is directed to integrated circuits, and more particularly, to a power down circuit for reducing power consumption in a
5 system-on-chip (SOC) comprising a plurality of circuit blocks by switching off the system clock to selected circuit blocks that are temporarily unnecessary.

Background of the Invention

Current trends in integrated circuit designs
10 call for creating an entire manufactured circuit system on a single chip. Such a system is termed system-on-chip or SOC. This differs from simple circuit integration in that many different types of circuits can be included on a single chip. For example, an SOC
15 could include a computer processor, various signal processors, a large amount of memory, various clocks, power down circuits, and necessary system controllers all integrated on a single piece of silicon or integrated into a single package. This level of
20 integration was not previously possible with prior integration techniques, and is very advantageous

10010728-110501
105011-3220100

because useful devices can be created in very small sizes.

Figure 1 is a block diagram showing an SOC 10a. The SOC 10a is formed of a number of different integrated circuit portions (IPs) or blocks 12, 14, ..., 20. Each IP block 12-20 is connected to a system clock 30. The system clock 30 distributes a system clock signal to each of the IP blocks 12-20.

Important examples of devices that can include SOC's are cellular phones, palmtops, notebooks, computer components, movable equipment, communication apparatuses, biomedical apparatuses, digital cameras, MP3 players, etc. Such applications generally require a battery or some sort of power supply, which typically presents cost, duration, weight and dimension issues. To increase the longevity of the power supplies for these devices, and especially for portable devices which require a portable power source, power consumption of the SOC's must be reduced from their current levels.

Dynamic power consumption of the different circuits blocks 12-20 integrated on a single SOC 10a is given by the formula $P=f \cdot C \cdot v^2$, where P is power, f is operating frequency of a circuit block, C is capacitance of all of the gates of the circuit block, and v is the power supply voltage. Therefore, in addition to reducing the power supply voltage and the overall capacitance, power of the SOC 10a may be conserved by reducing the operating frequency of the different circuit blocks 12-20. One way to implement this is to temporarily switch off the system clock for some of the IP blocks 12-20 of the SOC 10a that are not necessary for immediate functions. Because not all of

10010738-110501

the IP blocks 12-20 necessarily operate at the same time in the SOC 10a, some of them are unused and are eligible to be shutdown.

Figure 2 shows an SOC 10b that is similar to the SOC 10a of Figure 1, but additionally includes a power control manager 40. The power control manager 40 controls a bank of switches 42 that are coupled between the system clock 30 and the various IP blocks 12-20. When the power control manager 40 determines that particular IP blocks should be shutdown, such as circuit blocks 14 and 16, for example, a signal is generated and fed to the bank of switches 42. The bank of switches 42 then controls the particular switch coupled to the selected IP blocks, in this example IP blocks 14 and 16, and disconnects them from the system clock 30. When the selected IP blocks 14, 16 do not receive the system clock 30, they stop functioning and, based upon the above equation, do not draw any power because the operating frequency of the circuit is brought to zero.

Although the idea of separating the system clock from the various IP blocks is compelling, most SOC's cannot be controlled in such a manner. The implementation of such a system as shown in Figure 2 causes problems. As described above, many different types of IP blocks are contained within a particular SOC, and each of these IP blocks have unique requirements for when they can be safely shutdown.

It can therefore be difficult to establish an exact time when it is possible to switch off the clock to an IP block without causing errors. In some cases, if the clock to the IP block is stopped abruptly, there is a risk of preventing a critical operation of the

10010738-110501

block from being carried out. For example, an IP block could be performing a necessary communication protocol and the shutdown of the block could cause the SOC to disregard the protocol. Examples of protocols that could easily be disregarded include memory-DMA, and master-slave blocks among others. Additionally, removing a system clock from a counter or a timing signal generator could be fatal to that particular IP block.

- 10 Some of these problems are illustrated in Figure 3, which shows an SOC 10c that has prevented the system clock 30 from reaching the IP blocks 14, 16 and 18, while continuing to supply the IP blocks 12 and 20. In each of the cases of the non-supplied blocks 14, 16, 15 18, there are potential problems. For instance, the IP block 14 may be in the middle of a memory-DMA protocol operation with a memory unit 24, and its abrupt halt may violate that protocol. Similarly, the IP block 16 may be communicating with a slave peripheral 26, and an 20 abrupt halt may cause a malfunction or protocol violation. Additionally, the IP block 18 may include counters which rely on the system clock 30 for accuracy. Separating the system clock 30 from the IP block 18 could seriously degrade such accuracy.

25

Summary of the Invention

- In view of the foregoing background, an object of the invention is to accurately control the shutdown of multiple and different types of circuits blocks that are integrated into a single system to 30 preserve the necessary function of the circuit blocks.

This and other objects, advantages and features according to the invention are provided by

switching off the system clock for portions of the circuit blocks that are temporarily unnecessary. Specifically, this invention involves a power down circuit for use in a system-on-chip comprising a

5 plurality of circuit blocks each operating based upon a local clock signal. A system clock is coupled to one or more of the circuit blocks and provides a system clock signal that functions as the local clock signal of selected ones of the plurality of circuit blocks.

10 A power control manager is coupled to the plurality of circuit blocks and provides a signal that at least partially determines whether the respective system clock signals will function as the local clock signals for the corresponding plurality of circuit blocks.

15 More particularly, a communication protocol causes selected IP blocks to receive a shutdown signal from the power control manager. The selected IP blocks then complete their current activity and, on completion, switch off their internal clock and send an
20 acknowledging signal back to the power control manager. The shutdown signal is removed when the power control manager desires the IP blocks to restart, and the IP blocks send back an acknowledgment signal of the restart.

25 Based on this idea, this invention provides a selective power down circuit as previously indicated and defined in the characterizing portion of Claim 1.

Additionally, this invention provides a
30 method for powering down individual circuit blocks within a system-on-chip as previously indicated and defined in the characterizing portion of Claim 7.

1004073 110501

Brief description of the drawings

The features and advantages of the apparatus and method to power down selected circuit blocks within a system-on-chip according to the invention will be apparent by reading the following description of a preferred embodiment thereof, given by way of non-limiting examples with reference to the accompanying drawings:

Figure 1 is a block diagram of a system-on-chip according to the prior art;

Figure 2 is a block diagram of a system-on-chip that includes a power control management circuit according to the prior art;

Figure 3 is a block diagram of a system-on-chip highlighting the problems associated with the power control management circuit illustrated in Figure 2;

Figure 4 is a block diagram of a system-on-chip including the protocol according to the invention;

Figure 5 is a flowchart showing implementation of a first portion of the protocol according to the invention;

Figure 6 is a psuedocode listing describing operation of the flowchart illustrated in Figure 5;

Figure 7 is a flowchart showing implementation of a second portion of the protocol according to the invention;

Figure 8 is a psuedocode listing describing operation of the flowchart illustrated in Figure 7;

Figure 9 is a block diagram and a related timing diagram showing operation of selected signals within a system-on-chip including the protocol according to the invention; and

10010738-110501

Figure 10 is a block diagram showing implementation of portions of a complete system-on-chip including the protocol according to the invention.

Detailed Description of the Preferred Embodiments

5 Figure 4 illustrates interconnections that can be used to implement the protocol according to the invention. An SOC 100 includes a system clock 130, a power control manager 140, and two IP blocks 112 and 114. The system clock 130 is provided to each of the
10 IP blocks 112, 114. Additionally, two signal lines couple each IP block 112, 114 to the power control manager 140. The first of these is a power down request line 142, and the second is a power down acknowledgment line 144. Each IP block 112, 114 has
15 its own set of request and acknowledgment lines 142, 144 coupled to the power control manager 140. Of course any number of IP blocks 112, 114, etc. could be included in the SOC 100, with only the addition of the required number of request and acknowledgment lines
20 142, 144 and the proper connections to the system clock 130 being necessary.

 In operation, each IP block 112, 114 receives a "power down request" signal on the power down request line 142. A signal of either a 0 or a 1 is always
25 present on this request line 142. Normally, this signal will be a 0 when the IP blocks 112, 114 are in operation, but the 1 signal could be used instead, and such a change is within the scope of one skilled in the art. For purposes of this description, a 0 signal on
30 the power down request line 142 will indicate that the IP blocks 112, 114 should be operating normally, and a 1 signal on the power down request line 142 will

10010738-110501

indicate that the IP blocks 112, 114 should be shutdown.

When the power control manager 140 determines that a particular IP block should be shutdown, it puts
5 a 1 signal on the power down request line 142 coupled to the particular IP block. The selected IP block will receive the 1 signal on the request line 142 and finish its necessary operations. Once the operations are complete, the IP block will place a 1 signal on its
10 power down acknowledgment line 144. Placing this signal on the acknowledgment line 144 then causes the system clock 130 to disconnect from a local clock of the IP block, and the IP block stops drawing power.

Figures 5 and 6 illustrate a flowchart and
15 pseudocode, respectively, explaining operation of an implementation of the power down portion of the protocol. In Figure 5, the power control manager 140 desires the IP block 112 to stop drawing power, and issues a 1 on the power down request line 142. The IP
20 block 112 begins at step 200 and monitors the signal on the request line 142 in step 210. If the signal is a 0 in condition block 220, the IP block 112 continues looping through the steps 200, 210 and 220 until the signal on the request line 142 changes to a 1.

25 When the step 210 recognizes that the signal on the request line 142 has changed to a 1, it proceeds to step 230 where all of the necessary operations that are pending in the IP block 112 are completed. Once these operations are completed, the IP block 112
30 changes the signal on the power down acknowledgment line 144 from a 0 to a 1 in step 240, and ceases to function in step 250. The 1 signal on the acknowledgment line 144 is sensed by the power control

10010738-110501

manager 140. In the powered down state of step 250, the IP block 112 does not draw any power from the SOC 100. The psuedocode 190 of Figure 6 explains the above paragraph.

5 Figures 7 and 8 conversely show how the protocol operates as the IP block 112 is restarted after being shutdown. The IP block 112 begins in the state 250, the powered down state is also shown in Figure 5, and immediately reads the request line 142 in
10 step 260 and begins checking in step 270 to see if the signal on the request line 142 goes from a 1 to a 0. This indicates that the IP block 112 is to restart. Once the request line goes from a 1 to a 0, the system clock 130 (Figure 4) is again distributed to the IP
15 block 112 in step 280, and the IP block changes the signal on the acknowledgment line 144 from a 1 to a 0 in step 290. After this step, the IP block 112 proceeds back to the step 200, which is the normal operating step that the IP block started at in Figure
20 5. The psuedocode 194 of Figure 8 corresponds to the flowchart shown in Figure 7, and is self-explanatory.

The state of the request line 142 and the acknowledgment line 144 are stored in the power control manager 140. By evaluating the stored states, the
25 power control manager 140 can determine with certainty which state any given IP block is in, as illustrated in the following Table 1.

TABLE 1

142 Request Line	144 Ack Line	Status of IP block
0	0	Currently running

10010738-110501

1	0	Currently shutting down
1	1	Shutdown
0	1	Restarting

When both the request line 142 and the acknowledgment line 144 are both at a 0, the IP block would be operating normally. When the request line 142 goes to a 1 while the acknowledgment line 144 remains at a 0, that indicates that the IP block has just been instructed to shutdown, but is still finishing its required tasks before doing so. When both the request line 142 and the acknowledgment line 144 are at a 1, the IP block has shutdown and sends the acknowledgment of such back to the power control manager by placing a 1 on the acknowledgment line 144. Finally, when the request line 142 goes to a 0 while the acknowledgment line 144 remains at a 1, the IP block will restart operations.

Figure 9 shows a block diagram of an example IP block 112, along with a related timing diagram showing sample clock waveforms as they exist in the SOC 100 of Figure 4. Included within the IP block 112 of Figure 9 is a set of block logic 304, which is specific to the type of circuit contained within the IP block 112. Additionally within the IP block 112 is a shutdown circuit 300, which in one example can include a set of logic gates 306 and 308. In this particular embodiment of the shutdown circuit 300, the logic gate 306 is an AND gate and the logic gate 308 is a NAND gate, although any combination of logic gates that produce the correct result is acceptable for the

10010738 140504

shutdown circuit 300, and is within the scope of the invention.

In Figure 9, the NAND gate 308 has a first input connected to the request line 142, and a second input connected to the acknowledgment line 144. An output signal from the NAND gate 308 is a first input to the AND gate 306, with the system clock 130 being a second input. The output of the AND gate 306 is a local clock signal 310, which is fed to the block logic 304. As illustrated in Figure 9, the local clock will have the same frequency as the system clock 130, but will only be present when the output signal from the NAND gate 308 is a 1 signal.

Examples of signals feeding the shutdown circuit 300 are also shown in Figure 9 for three different time periods t1, t2 and t3. In all of the time periods t1, t2 and t3, the system clock 130 continues to operate at the system frequency. In the first time period t1, the request line 142 changes from a 0 to a 1. This indicates that the power management system 140 of Figure 4 desires the IP block 112 to be shutdown. The IP block 112 begins to shutdown at the end of the period t1, which correlates with the step 230 shown in Figure 5. Because the acknowledgment line 144 is still set to a 0 throughout the entire period t1, the local clock 310 would continue to be supplied to the block logic 304 during the entire period t1.

In the period t2, the IP block 112 completes its current operations and raises the acknowledgment line 144 from a 0 to a 1. Once this occurs, the output of the NAND gate 308 goes LOW, and therefore the output of the AND gate 306 also goes LOW. This causes the

10010738-110501

local clock 310 to stop, and the IP block 112 is in a powered down mode.

In the period t_3 , the request line 142 changes from a 1 to a 0, indicating that the power control block 140 desires that the IP block 112 restart its operations. When the signal on the request line 142 changes from a 1 to a 0, the output of the NAND gate 308 immediately (after a negligible propagation delay) changes from a 0 to a 1. This, in turn, causes the AND gate 306 to again pass the system clock 130 as its output for the local clock 310, which is again fed to the block logic 304. Once the local clock 310 is present within the block logic 304, the IP block 112 lowers the acknowledgment line 144 from a 1 to a 0, indicating that it has resumed operation.

Figure 10 shows a top level architecture implementation of the protocol according to the invention. An SOC 400 includes IP blocks 412 and 414. Again, any number of IP blocks could be present within the SOC 400, and only two are shown for purposes of illustration. A system clock 430 is always in operation within the SOC 400, and is distributed as a first input to an AND gate 406 within each of the IP blocks 412, 414. Another input to the AND gate 406 is an output from a NAND gate 408, also present in each IP block. The NAND gate 408 has a first input from a power down request line 442, and a second input from a power down acknowledgment line 444. When the signals on the request line 442 and the acknowledgment line 444 are both 1's, a local clock 410 is not passed to the respective block logic within the IP blocks 412, 414. Otherwise, the local clock 410 is the same as the system clock 430, as discussed above.

10010738 110501

The power control manager 440 includes a set of two registers, a first register 446 and a second register 448. These registers each contain memory storage locations, at least one location for each IP block 412, 414 within the SOC 400. The first register 446 is coupled to all of the power down request lines 442 in the entire SOC 400. That is, each of the power down request lines 442 will have a 0 or a 1 signal on it determined by the data stored in the respective memory location of the first register 446. Providing data on a signal line, such as the request line 442 to match data stored in a memory location, and reading data from a signal line and storing it in a memory location are conventionally known.

In one embodiment, a CPU 450 can write data into the particular memory location of the first register 446 for a particular IP block within the SOC 400, and the power down request line 442 will be changed accordingly. In another embodiment, the CPU 450 would not be allowed to write data into the first register 446, but could only read data already written there by the power control manager 450. In still another embodiment, programmable control would be given where it could be selected whether the power control manager 440 or the CPU 450, or both, could write data into the first register, thereby controlling the shutdown of the relative IP block.

The second register 448 is coupled to all of the power down acknowledgment lines 444 in the entire SOC 400. Each of the power down acknowledgment lines 444 will have a 0 or a 1 signal on it determined by the signal placed on the acknowledgment line 444 by the respective IP block 412, 414. Because only the IP

block itself can change the signal on the acknowledgment line 444, neither the power control manager 440 or the CPU 450 can write data into the second register 448, but both of them can read the data stored there.

10 An advantage to implementing the inventive protocol in the manner shown in Figure 10 is that the power control manager 440 and the CPU 450 always knows the current states of the IP blocks 412, 414 in the SOC 400 by comparing the data stored in the particular locations of the first and second registers 446, 448 that denote the respective IP blocks, and comparing the data read from the registers to the table provided above.

15 This protocol provides a straightforward and convenient way to safely switch off the clock to desired circuits within a system-on-chip by providing a signal to the desired circuits and letting them finish their processing prior to being shut down. The
20 implementation described above provides a further benefit in that control of such shutdowns can be executed by hardware and/or by software.

1010728-110501